

White Paper Report

Report ID: 98504

Application Number: HD5063309

Project Director: Paul Kaiser (paul@openendedgroup.com)

Institution: Unaffiliated Independent Scholar

Reporting Period: 4/1/2009-4/30/2010

Report Due: 9/30/2010

Date Submitted: 9/29/2010

Spatializing Photographic Archives

White paper for the National Endowment for the Humanities
Digital Humanities Start-up Grant HD-50633-09

Marc Downie and Paul Kaiser — September 30, 2010

For this project we have built an open-source tool that exploits recent advances in computer vision to put these advances within reach of the scholar.

Specifically, we have refined an algorithm that generates 3D datasets from unstructured collections of photographs. We have then created an open-source tool (the *Spatialization workbench*) that permits scholars to create such reconstructions and that gives them new ways of examining the resulting 3D “point-cloud,” with every photograph and the corresponding position of the photographer situated in the three-dimensional space. These visualizations may be inspected in stereoscopic 3D display.

The Spatialization workbench is built within our larger open-source software platform, *Field*, which can create dual interfaces: one for novice use of a graphic user interface as is; and one for expert use at the coding level to freely modify both that interface and its underlying capabilities. The latter mode allows scholars to precisely tailor the software to the precise needs of their research domain.

Our hope is that the value of the workbench software goes beyond providing scholars with a useful new method of research, but also serves as a model for a transformative open-ended approach to the design of scholarly tools.

Our further goal has been to apply the spatialization technique to examine photography in a radical new way by stepping outside its two key defining properties: the framing of a subject and the freezing of a moment. Our main test case is noted American photographer Richard Misrach’s landscape photography.

We have reconstructed two key sites of photographs Misrach took in the mid-1980s, placing his original photographs within our point-cloud reconstructions. We have also shown how to record what we

call the choreography of framing by reconstructing Misrach’s path and composition as he photographed a new subject for a work-in-progress.

Using metadata from a GPS equipped camera, we have tied photos to online geospatial databases (Google Maps and OpenStreetView).

We have also applied our method to the much more constricted possibilities of an older archive containing surveillance and combat photography from the battle of Okinawa in 1945.

We provide a proof-of-concept for a simple capture pipeline that would allow scholars to capture an object using the everyday technology of smartphone and laptop computer. We hope to suggest a new form of 3D notebook for scholarly fieldwork, analysis, interchange, and presentation, and so we have contrasted the ad hoc and inexpensive approach this would allow with the more expensive and demanding method of laser range-finding capture.

Our tools are already scalable and can be deployed on ad hoc networks of Macintosh computers. We discuss an imminent expansion of this scalability funded by the National Science Foundation, which will enable the data-mining and visualization of large data-sets, then give two hypothetical examples of how this might be applied to much larger photo/video reconstructions.

We end by suggesting a new form of interdisciplinary project — one that requires a collaborative team with expertise in oral and cultural history, sociology, ecology, computer science, photography, and immersive installation art. We argue that this new form of visual scholarship would be intrinsically hybrid — an artwork that enfolds many art expressions; a history that intertwines personal, cultural, historical, and ecological strands.

NOTE ON ILLUSTRATIONS

This report tells its story as much in pictures as in words. But since our subject is the computation of spatial depth from photographs, the flat page presents us, and you, with a challenge.

2D images on the page lack two of the key means we have for seeing depth. They allow neither *stereopsis* (for the flat page presents each eye with an identical image) nor *parallax* (for when you move your head, that image doesn't shift).

In our actual workflow, we use active-stereo 3D, which produces a near-perfect illusion of depth with no color distortion — and the illustrations come alive.

While this report gives you access to the necessary information and code to eventually duplicate this projector-based, active-stereo set-up yourself, we've sought a way for you in the meantime to get a small sense of how much better all this works in 3D than on the page.

To that end, we've provided 3D animations of several of the project's test spatializations, which are referenced in the relevant sections below. For now, we've had to make do with the old and largely outmoded anaglyph technique (popularized in the 1950s), which has you wearing glasses with red and cyan filters to view what will then appear to be grayscale images floating in 3D space. Not great, but good enough.

The advantage of this approach is its simplicity and affordability. An internet search will turn up numerous sources of anaglyph glasses — at the time of this writing, a paper pair could be had for 99¢.

TABLE OF CONTENTS

Abstract	i	<i>Desert Fire</i>	33
Note on illustrations	ii	The Thousand Palms site	34
Overview	1	The spatialization	35
Algorithm illustrated	2	Choreography of framing	39
Open v. closed	9	Bulldozer spatialization	42
A case in point	8	Integration with other data: GPS/mapping	44
An alternative: Field	10	Other forays	47
New spatialization workbench	12	Battlefield archive — Okinawa, 1945	48
The interface	13	Wartime photography spatialization	49
Malleability	15	Towards a scholar's 3D notebook	52
Richard Misrach's landscape photography	17	Antique clock spatialization	53
Landscape photography of the American west	17	Reflections	54
Rephotography — the act of framing	17	Limitations	54
<i>The Desert Cantos</i>	20	Field capabilities & expansion	56
<i>Submerged Trailer Home</i>	21	Advantages of scalability	56
The Bombay Beach site	22	Hypothetical example of data-mining — Ground Zero	58
The spatialization	23	Hypothetical example of interdisciplinarity —	
Continuous change	32	Bombay Beach	59
		Acknowledgements — People	63
		Acknowledgements — Open-source projects.....	64

Computer vision

Seeing the world by looking at first seems effortless to humans — not so to computers.

In 1966, MIT's Marvin Minsky faced the problem of robots navigating a space and asked an undergrad to take the summer to solve it. That effort, which spawned the field of computer vision, is still going strong 40 years later.

One of the perennial problems has been how to build up a perception of three-dimensional objects from 2D views of those objects. This *structure from motion* problem is typically seen as one of the key abilities needed for a robot to avoid obstacles and reach goals.


Taking a different tack in **Photo Tourism: Exploring image collections in 3D** (ACM Transactions on Graphics: Proceedings of SIGGRAPH 2006), Snavely, Seitz, and Szelisk realized that you could calculate 3D shapes and build up 3D spatializations from unstructured collections of photographs. Rather than being captured from a robot systematically moving around the world, such photos could be downloaded from diverse collections found here and there on the Internet.

Ours is a time that has seen ever more information recorded photographically. The ability to snap a picture or to capture a movie is at hand's reach: it's in your pocket (as your cellphone or Flip camera) or in your bag (as your laptop's video conferencing camera). And while relatively inexpensive, these devices are increasingly capable, with high-definition video and even high-dynamic-range photography quickly becoming standard. Meanwhile, hard drives have kept growing in size and shrinking in price, while Internet storage and display (Flickr, Picasa) has become free. And so the cost of taking, storing, and sharing photographs — in terms both of money and of effort — has dropped to nearly nil.

Online, the resulting flood of pictures taken by individuals has joined those gushing from government agencies (NASA's satellite photography) and private companies (Google's StreetView). With the world blanketed with photographs, undreamed-of possibilities have opened up for humanities scholars. This project explores one such opportunity that has emerged from the field of computer vision. We have built an open-source tool that exploits recent advances in computer vision and can put these advances within reach of the scholar.

Methods have emerged that can extract three-dimensional information from an unstructured, un-annotated collection of related photographs. Specifically, given multiple shots of the same scene, one can now (1) recover the location and angle, in 3D, of the camera that took each shot and (2) reconstruct, in 3D, much of the geometry of the original location. With minimal human intervention, hundreds and thousands of images can be projected back into a shared three-dimensional space. This can yield astonishing results: it is as if one can now peer around the constricting frames of photographs, seeing more there than one ever thought that flat photographs could reveal.

We illustrate the basic technique by walking you through it on the next six pages.





1. Image collection

Photograph a given scene from as many different angles and distances as you can practically manage. Your approach can be unsystematic so long as you take a wide variety of partially overlapping shots. (These are drawn from the Bombay Beach site discussed on pp 21–32.)




2. Feature detection.

The software analyzes each image to find its potentially significant features/landmarks. These are indicated by the red crosses above.




3. Feature matching.

Now the software computes a large dictionary of possibly common points, then starts to match likely pairs. Above are pictured two shots of the same scene, with connecting lines showing the best points in common.



4. Spatial alignment of paired photos.

Given corresponding points, the software is capable of discovering two things at once: (1) the positions and lens parameters of the cameras; and (2) the 3D structure of the feature/landmarks in the scene itself. This allows the photograph to be aligned and the 3D “point-cloud” representation of the scene to be created.




5. Multi-camera optimization.

By cleverly optimizing multiple images simultaneously, the software can reconstruct wide expanses of 3D space from a large set of images. While we only show 2 images here, the red dots represent the positions of reconstructed cameras. (This scene is constructed from roughly 400 photographs.)



6. Alignment of actual and virtual cameras.

By placing the virtual camera at the position from which the photograph was taken, and matching focal lengths, we can look through the photograph into the virtual scene.



7. Further analysis.

The point-cloud can then be subjected to further analysis, and the software can begin to infer surfaces and edges, as rendered by our software. This may be seen in anaglyph form at <http://openendedgroup.com/index.php/software/spatial/bb1-anaglyph/>.

Many of the most interesting computer science research projects languish in academia. Going no further than a concept paper and a proof-of-concept implementation, they frequently find themselves orphaned. The common practice of releasing the underlying research code as open source, a fine thing in principle, is often ineffective in practice — for usually such code, however brilliant, remains too specialized to be used by others outside the original group and fails to find its way out into the wider world.

It is unreasonable to expect research engineers to do anything different. Their business is not to produce complete software packages for general use, but rather to demonstrate that their algorithms work, and to do so in the most expeditious manner possible before moving on to their next problem. Thus they feel no compunction about using obscure file formats, providing minimal interfaces, and exploiting back channels for code distribution — whatever works.

But if we can't expect engineers to provide us with intelligent, usable interfaces for far-reaching algorithms, nor should we expect commercial software companies to do so. Even when they package up a new bit of cutting-edge research, they typically do it in a closed and inflexible form. While this may well serve the minimal needs of the general public (their customer base), it will fall well short of what a scholar requires for advanced research.

A case in point

The story of the basic algorithm developed at the University of Washington by Slavely, Seitz, and Szeliski — which forms the basis for the spatialization of unstructured collections of photographs — clearly illustrates what can happen.

Released as a command-line program in Linux, with no graphical front-end, their open-source code had a limited intended audience, though it provided the basis for their remarkable SIGGRAPH paper of 2006.

The Microsoft legal scarecrow

The terms of use for the Photosynth website state:

You will not disassemble, decompile, or reverse engineer any software included in the service, except and only to the extent that the law expressly permits this activity.

But what does this mean exactly? The second clause puts the bluster of the first in doubt — as well it should, since the key underlying code is open-source and thus precisely intended (and legally protected) for re-use and re-engineering.

Microsoft then stepped in to provide a graphical front-end for this 3D reconstruction technique, releasing it as a closed and proprietary product they dubbed “Photosynth.” This now functions as a website attraction — www.photosynth.net — and has become a vehicle for virtual tourism. Users can upload their sets of photographs, which are then automatically processed by the spatialization algorithm and, if successful, presented online. The site has thus become a large repository for virtual walkthroughs of such destinations as Yosemite, San Francisco, the Sphinx, and the Roman Forum.

While well-designed in its way, the site is more toy than tool. It’s been so carefully crafted with an eye towards ease-of-use that its *range* of use is severely constricted. (Imagine a word-processor that allowed the entry of text, but not the editing, printing, and sharing of it.)

It’s a one-way street: you get no data back from Microsoft. Thus you can create none of the measurements, analyses, and visualizations of the data that would be truly useful for serious investigation. You can’t back out the exact position and angle from which a given photograph was taken; you can’t align one scene with another; and you can’t connect a scene to other relevant data sources (for example, mapping data). No good: scholars should own their own notes.

An alternative: Field

So what would serve the scholar better? It won’t do to simply create an open-source duplicate, or even an improved version, of the Photosynth software if it merely recapitulated the closed and finished commercial software product (too often the strategy of the most widely disseminated open-source packages).

No, what’s needed is something different: a platform for experimentation and scholarship. Not a carefully designed and controlled finished interface for the underlying algorithm, but an

History of Field

The development of Field began in 2001 when Downie was a doctoral student at MIT's Media Lab. The codebase he began developing played a key role in interdisciplinary research that combined work in animation, biology, computer science, and robotics.

Following graduation, Downie continued to develop this codebase as a member of the OpenEnded Group, a digital artists collective that creates large interdisciplinary art-works drawing upon such fields as dance, music, data-mining, lighting design, live cinema, artificial intelligence, and public art.

For NSF-funded work at Arizona State University, Downie reoriented his codebase to enable the convergence of dance, interaction, live motion capture and analysis, and live digital imagery. The newly organized code, now given its present name of Field, allowed the visual artists to iterate as rapidly with their imagery as their collaborator, choreographer Trisha Brown, could do with her dancers.

In 2007, Field became an open-source project with seed funding provided by the Mellon Foundation. This was to support the *Loops Preservation Project*, which innovated ways to preserve both choreographer Merce Cunningham's signature solo dance *Loops* as well as the advanced digital artwork deriving from it, also entitled *Loops*.

In 2008, Downie published an overview of the software in **Field—a new environment for making digital art** (*ACM Computers in Entertainment*, vol. 1, issue 4).

In 2009-10, Field received related support from two further sources. Portland Green Cultural Projects (with funding from the UK Arts Council) underwrote the development of the Choreographic Language Agent, a tool for experimental choreography, the creation of which led to a further refinement of Field — the addition of a multi-modal interface to accommodate expert and novice users.

In the same period, EMPAC (Experimental Media and Performing Arts Center) of Rensselaer Polytechnic Institute provided indirect support for several extensions to Field, most notably the addition of music analysis and high-definition 3D — the latter especially useful for visualization, which is made far more powerful by the addition of depth and the enlargement of scale.

environment for creating, editing, and ultimately sharing such interfaces — ones tailored to the specific needs of a scholar's investigation.

We've developed an open-source software platform for interdisciplinary research called Field that embodies this principle. Field is what you might call a *meta*-authoring system, for it allows you to modify and create your own interfaces for whatever problem domain you wish to explore.

Since we've used Field to author the photo spatialization software for this project, a quick overview of Field is called for here.

The founding design principle of Field is that large quantities of complex code written in a safe, robust language (Java) are rapidly tied together with small amounts of simple code written in a lightweight, dynamic scripting language (by default, Python). The Field architecture comprises a minimal core to which is added a powerful plug-in system that can call upon domain-specific tools, libraries, and languages.

Field's all-embracing approach to other software systems stems from the realization that since the best work available in traditional domain-specific disciplines is already coded, such code is better appropriated than re-engineered. Thus Field can easily wrap up, examine and choreograph other programs and libraries. (For this project, Field has either built upon or fully incorporated the elements listed on p. 64.)

Field provides a hybrid text editor that can incorporate text-based programming at the same time as graphical user-interface elements. Its flexible and extensible interface gives space as necessary for notation, visualization, interface construction, and debugging. In addition, Field makes available an advanced GPU-accelerated graphics system.

The spatialization workbench is implemented in Field, the software development environment described above. The workbench is what we call a *sheet* in Field — a sheet being a work area in which you both write code interactively *and* interact with the results generated by that code.

We’ve already constructed a spatialization workbench sheet for you, so it’s ready to use as is. It does two principal things:

1. It spatializes any unstructured sets of overlapping photographs that you load into it.

In technical terms, it does this by providing a robust, restartable wrapper around the spatial reconstruction algorithm developed at the University of Washington, to which it adds these refinements:

- ground plane detection
- surface-normal estimation
- shape-finding (planes)
- point outlier removal
- point-cloud densifier.

2. It displays the point-clouds created in the spatialization process, which you may view the spatialized scene from any vantage point. You may also see the position and angle from which each original photo was shot, and examine the photo itself in its precise spatial context.

In creating the workbench sheet, our goal has been to shield you at first from the underlying complexities of its code, but not to close them off to you for good. *To the contrary*: you may modify the software to suit the needs of whatever project you have at hand. Field’s underlying code is open to examination and ad hoc modification, as the simple example on p. 16 illustrates.


Technical background & requirements

Since the software has been compiled for recent versions of Apple’s OS10 rather than for Linux, it has no outside dependencies, and installation is a straightforward process that results in a simple double-clickable application. It requires a recent model of Macintosh — for larger spatializations, a MacPro.

To download the program, and for further documentation and support, see: www.openendedgroup.com/field/wiki/ReconstructionDistribution.


The workbench, as well the larger Field platform on which it depends, are open source, released under the community standard version 3 of the GNU Public License.

The interface



- ① **CREATE** — to select a set of photographs for spatialization.
- ② **SPATIALIZED SETS** — lists processed sets with current status of processing. Right-clicking loads a set into the viewer.
- ③ **LOAD** — to load a spatialized set into the viewer for sets not listed in ②.
- ④ **VIEWER** — displays the spatialized point-cloud you loaded. You can pan, zoom, and rotate the view with your mouse.


- A red dot in the displayed point-cloud marks the location of each original camera — mouse-clicking the dot reveals that camera's frustum and photo; double-clicking it moves the current view to match that of the photo. (see process illustrated on the next page.)
- ⑤ **POINT SCALE SLIDER** — to adjust the size of points in the point-cloud. (see p. 15)
 - ⑥ **RESET VIEW** — to reset display view to match that of the first photo in the set (also initial first view on loading).




Double-clicking a red dot — which indicates the position of an original camera — pans, rotates, zooms, and moves the display view to align it with the camera's. Above: 3 successive re-alignments of view.

Malleability

To illustrate the ways in which you can modify the workbench both at the graphic interface and at the code levels, let us examine adjustments you can make to the point-cloud display.



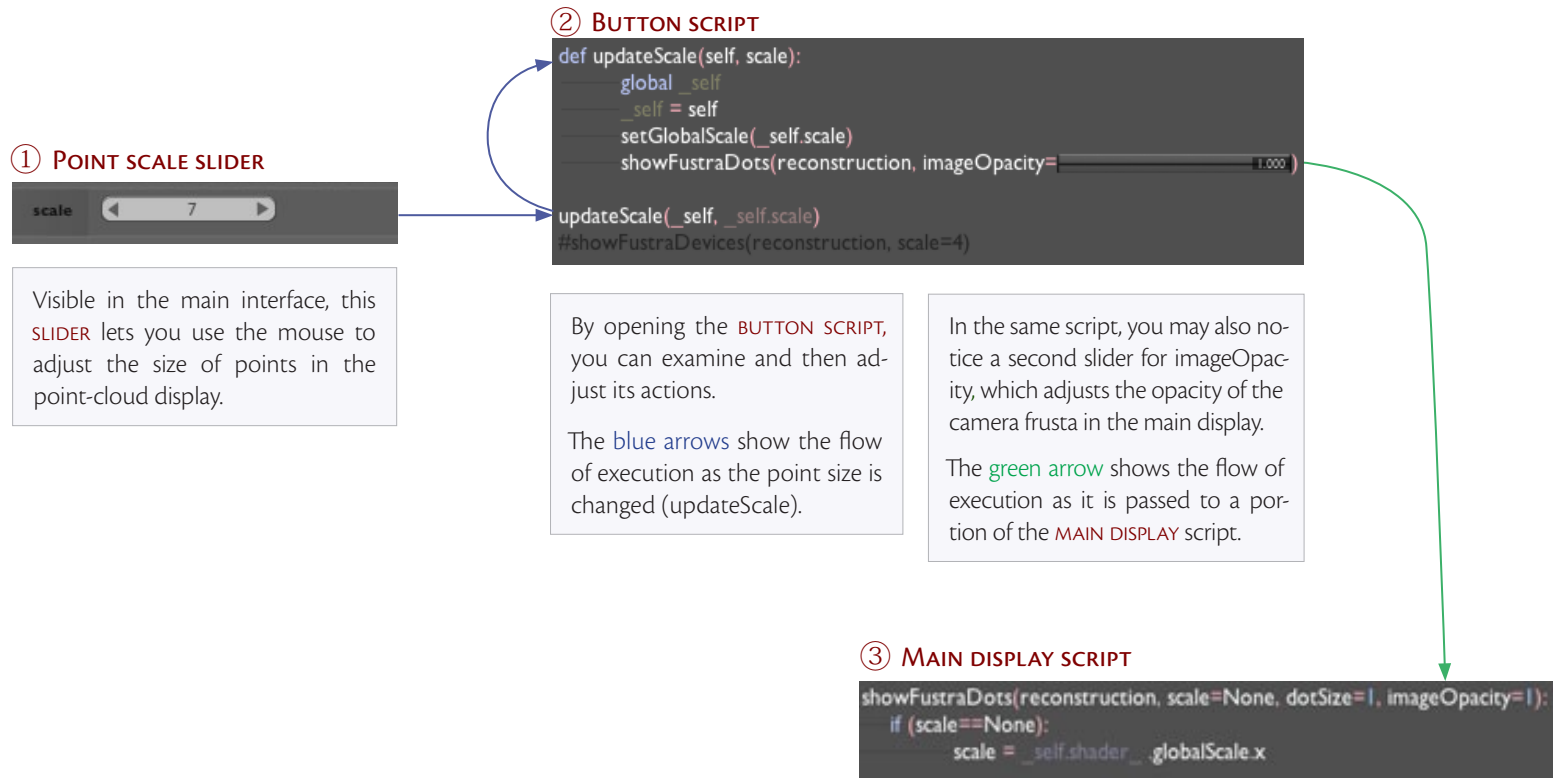
With the mouse, you can make the points making up a point-cloud smaller (*left*) or larger (*right*).



In code you can do much more. For example, you can make the photograph displayed in a camera's frustum fully opaque (*left*) or partly transparent (*right*).

The software lets you examine — and change — its underlying code and functionality. Since Field’s innovative text-editor also features several graphic-user-interface elements, simpler adjustments can be surprisingly straightforward.

Here’s how the two adjustments illustrated on the previous page work.



The advance of technology in the American West

Rebecca Solnit has shown how tightly coupled were the exploration and industrial conquest of the American West with advances in photographic technology, which led via Eadward Muybridge to the invention of the moving picture (and in turn to the defining genre of the Hollywood Western). See Solnit's *River of Shadows: Eadward Muybridge and the Technological Wild West* (Penguin: 2004).

We note how appropriate a proving ground is the Western landscape to test advances in photographic survey technologies.

Landscape photographers frame and then freeze what they encounter on the land, snatching an instant from the flux of time and of change.

But a landscape photograph is necessarily a recording of the world, and not just the intended and foreseen recording — however constricted its frame, artful its composition, hidden its bias, or artistic its intention, unexpected traces of the world imprint themselves on the image in the moment of exposure. These traces often remain invisible to eye and to mind until later examination (it's tempting to say *cross-examination*) in a different context, sometimes years or even generations afterwards.

Nowhere has the re-examination of landscape and photography taken a more interesting turn than in the American West.

In the 19th century, photographers accompanied the earliest geologic and geographic surveys of the region, which they framed and defined to mold cultural conceptions of the West that persist to this day. These photographs surveyed the western reaches of our continent not only to map their riches for possible exploitation, but also, as Joann Verburg has written

to let people who couldn't go there know what it was like. In fact, these photographers were following their own artistic visions. They photographed views considered beautiful by the aesthetic standards of the day. In the process, most of their work tamed the West as it brought home to easterners pictures of places where few of them could have survived.

Rephotography — the act of framing

Verburg drew these conclusions after participating in the Rephotographic Survey Project, an effort led by the photographer Mark Klett from 1977 to 1979. The project tracked down the sites and precise vantage points from which more than 120 photographs of the original 19th century surveys had been shot. Matching as

Rephotographic Survey Project books

The Project's first book is a tour-de-force. **Second view: the Rephotographic Survey Project** (Klett, Manchester, Verburg; University of New Mexico Press, 1990) unites theory, method, and execution to produce flawlessly paired photographs a century or more apart.

Klett's follow-up work two decades resulted in **Third views, second sights: a rephotographic survey of the American West** (Museum of New Mexico Press in association with the Center for American Places, 2004).

A useful companion volume is **View Finder: Mark Klett, Photography, and the Reinvention of Landscape** (University of New Mexico Press: 2001). Written by William L. Fox, one of Klett's collaborators on *Third Views*, it combines detailed accounts of Klett working in the field with historical and conceptual reflections.

closely as possible not only the framing of the original shots but also the light conditions, the team rephotographed the same sites as they presented themselves to the lens 100 years later.

The stunning success of their project derives from the rigor of its execution. For example, before rephotographing a Timothy O'Sullivan shot of the Canyon de Chelly, Klett first had to figure out that its precise lighting conditions could be duplicated only in a daily three-minute interval during a two-week window in August.

It was this sort of fanatical precision that makes the project's pairing of views so authoritative and telling, for not only do shadows align precisely—even cloud formations and water reflections are barely distinguishable from original to rephotograph. This precise visual matching neutralizes the aesthetic difference between the shots, so that it is only differences in the actual landscape that register.

Such differences sometimes defy prediction. Almost as often as the later view shows a subsequent man-made alteration of the landscape, it instead reveals the complete erasure of an earlier one — in one instance, the near-complete disappearance of a quartz mill from the landscape.

However ambiguously, the pairing of photographs brings time back into play, forcing the original to be re-examined as part of a sequence. Klett observes that

The earliest pictures mark the starting point, but no one image, first or latest in the series, represents a definitive statement. Each is simply another perspective in an ongoing time-image network. The individual pictures are like single frames picked from a time-lapsed film; when they are viewed in succession they give the appearance of time in motion, of continuous change. However, in this case, they are perhaps the first two frames in a monumental film, one which will span thousands of years.

In follow-up work two decades later, Klett sought to expand the scope of the rephotographic project. Revisiting and rephoto-

graphing many of the same sites as in the original effort, he tried finding ways to convey what lay outside the rigid compositional frame, not only visually but aurally and contextually. Thus, in the work presented in his combined book/CD-ROM publication *Third views, second sights*, he included not only the paired photographs, but also videos and snapshots of the sites and surrounding areas; diary entries and field notes; audio recordings of interviews, ambient sounds, and found sound artifacts (discarded LPs, cassettes, etc); artifacts and objects found at the sites; and Quicktime VR panoramas. One of William Fox's field notes, describing the interior of a ruined house at the site of another Timothy O'Sullivan photograph, gives a sense of what they were after:

The sense of time periods folding into one another multiplies as Neolithic chert flakes turn up next to shiny brass 30.06 rifle casings. An exquisitely eroded book found inside the house has "Our Nation at Dead Center" at the top of the verso page; "Exercises and Activities" reads the recto. Byron holds up a square, wafer-thin record to the sky, the grooves of the transparent text of religious homilies shot through with sunlight. Layers of sequential habitation are visible here, all connected in a single space yet mysterious as to purpose.

| 19

This sense of layered time and space is well worth pursuing and pinning down, but Klett's second rephotographic effort fell well short of his first, for a reason most instructive for our purposes here.

What focused his first project was its perfect formal framework. Lacking this, the second project had no structure by which to order the sprawling contextual materials it gathered, thus falling short of its promise.

Would that promise be borne out better given the right formal framework? More to the point, might the spatialization of photographs provide just that?

To explore these possibilities, we investigated not the archives of 19th century photographers, but rather that of contemporary photographer Richard Misrach.

Related books by & about Richard Misrach

Though Misrach's *Desert Cantos*, as well as his other works, are best seen in the large-scale prints he exhibits in museums and galleries, his **Chronologies** (D.A.P./Fraenkel Gallery, 2005) is so big and so lavishly printed a book that it provides a near-substitute experience.

More affordable and readily available are **Desert Cantos** (University of New Mexico Press, 1987), which includes photographs from *The Flood* and *The Fires* cantos, and **Crimes and Splendors: The Desert Cantos of Richard Misrach** (Tucker, Misrach, Solnit; Bulfinch Press, 1996).

Two further volumes — **Bravo 20: the bombing of the American West** (R. Misrach, M.W. Misrach; Johns Hopkins University Press, 1990) and **Violent Legacies: Three Cantos** (Sontag, Misrach; Aperture, 1994) — consolidate Misrach's most visceral yet conceptual readings of the western landscape.

The Desert Cantos

Richard Misrach is one (and Klett another) of a generation of photographers who challenged earlier renditions of the American West. From the start, photographers of the western landscape had isolated and exalted the purity of nature. Training their lenses on pristine expanses and distant peaks, earlier photographers tended to hide or reduce human presence in their idealizing compositions, as a glance at Ansel Adams' fine art prints or any Sierra Club calendar will attest — or as shown too, even now, in the pictures taken by the majority of serious amateurs.

In 1979 when Misrach started his ongoing lifework, the *Desert Cantos*, he was intent on putting back in everything his predecessors had left out. The *Cantos* are, in his words,

an attempt to provide an alternative, more accurate way of understanding the West. They show a land not of open spaces and wilderness -- of loners subsisting on the earth's natural bounty, and heroic efforts to "civilize" the West — but a land used by military and government agencies for the development of mass destruction, a land where natural resources are poisoned.

His subjects have ranged from an Air Force bombing range in Nevada to car racing on the Salt Flats of Utah to dead domestic animals piled in pits in the desert. He doesn't turn away from the beauty of the desert landscape — instead he finds that beauty not only in its skies and dunes and horizons, but even in its most terrifying disfigurements.

You will not find in Misrach many instances of Cartier-Bresson-like "decisive moments" caught on the fly. Much more often his pictures come only after his patient observation of a site, often periodic and at times stretching over the course of years. He waits until all the elements have converged in his frame perfectly — light and shadow, and the minute details of land and sky. This happens for him most often at dusk or at dawn, in the light he favors.

It was two of Misrach's sites that we took as our testing grounds.

Submerged Trailer Home

Misrach shot *Submerged Trailer Home* in Bombay Beach, a very small town of semi-permanent abodes arranged in a grid of streets 5 wide by 9 long on the western shore of the Salton Sea in southern California. Itself the product of a truly vast water engineering accident in 1906–08, the Salton Sea had overflowed its banks in the mid-1980's, swamping many of the dwellings clustered on its beaches.

A quintessential Misrach location, the Salton Sea fuses beauty with catastrophe. It boasts at once the most polluted and noxious of waters and the most magnificent formations of white pelicans overhead. Its habitations may be the stranded vestiges of a time when the Sea was promoted as an American Riviera, a place of fishing and water-skiing and clean desert air — but now, in the isolation brought on by their ruin, these spots can still entrance you with views of sunsets and starscapes, of dappled expanses of blue water and the delicate outline of mountains on three of its horizons.

You can approach the Salton Sea as William Vollman does in this recent passage:



When each shore is a far shore, when Imperial defines itself gradually through its long boxcars, hills, palm orchards, vineyards, and the blue pallor of the Salton Sea beyond, then the pseudo-Mediterranean look of the west side as seen from the east side (rugged blue mountains, birds, a few boats) shimmers into full believability. Come closer, and a metallic taste sometimes alights upon your stinging lips. Stay awhile, and you might win a sore throat, an aching compression of the chest as if from smog, or honest nausea.

Our approach to this site, proposed to us by Misrach as a practical object for our investigation, was disbelieving at first. Gazing at the photograph he'd sent us, we thought it the least likely candidate for spatialization. We knew that the waters had long since receded there, so what could possibly be reconstructed from the ruins they had left behind?



SUBMERGED TRAILER HOME, SALTON SEA, CALIFORNIA — MISRACH 1985

Vollman chronicles the twisted history and blighted beauty of the Imperial Valley region (expansively defined) in his recent book *Imperial* (Viking Adult: 2009), a pungent and exhaustive tome.



"rugged blue mountains, birds, a few boats" — contemporary Salton Sea tourist postcards.



THE BOMBAY BEACH SITE

The site we investigated is pictured in the panorama above, stitched together from multiple snapshots taken from the raised berm that borders the roughly rectangular area on two of its sides.

It was the existence of this berm that originally protected the rest of Bombay Beach from the 1980s flooding — and that gave Misrach the dry vantage-points from which he could frame *Submerged Trailer Home* and several other photographs.




He'd brought workprints of these to the site in October 2009 to show Kaiser what half-ruined landmarks he could still match up to his photographs. There were not many of these shots since in

those days Misrach used a heavy 8x10" Deardorff camera and carried a maximum of 25 double-sided film holders. Limited in the number of shots he could take in a day, he was parsimonious in doing so.


This meant that there were now not enough older photographs for us to respatialize the original space, so we resorted to a different strategy — rephotographing the entire space and then setting two of the original photographs in the context of that new spatialization.

On two successive days Kaiser shot the site from a variety of angles and distances, working only at mid-day to keep the shifting of shadows to a minimum. In rapid-fire mode, he ended up taking more than 8,000 shots, enough for us to spatialize the entire site and enabling us to match up key remaining landmarks with those visible in the original two photographs.




Spatialization of Bombay Beach site

A successful spatialization allows you to view it from any angle, so here we show the site from bird's eye view, numbering the landmarks that we could match to Misrach's two original photographs. (How much more legible you would find this is in 3D!)



Matching of structures then & now

We matched three ruins to structures visible in the *Submerged Trailer Home* photograph, as detailed in the following pages. Of the trailer home itself, nothing conclusive remained, though where it once stood we found a mess of debris.




① **The doghouse.** The structure's shape, the incised crescent moon, the little smokestack on the roof — these distinctive features made this the easiest match, though the structure listed to its side atop its ruined foundation.

"Doghouse" was our tentative designation, a mere guess at its actual function, which, come to think of it, must have been different: for where was the door?




② **Truck shell.** We were surprised to match the white truck in the original photo to a ruined truck shell amid the brushwood. It was covered with spray-painted graffiti both outside and in — various initials of the *I was here* sort. As for the original red truck, it had left no trace.






③ **STAIRWAY HOUSE.** This house still had its distinctive outer staircase, though the rooftop deck to which it led had collapsed completely. One could still venture to climb up the creaky steps, the top one affording a useful raised vantage point from which to take downward shots of the site from angles different from those to be had on the berm.

A large sheet of corrugated metal was propped against the back of the house, it too covered with spray-painted initials.





A second unpublished Misrach workprint from the same day took in the same scene from a wider angle, revealing two more structures still evident on the site 20 years later. The most imposing one was:

④ The Brad house — so named because we found that message spray-painted on its rear wall (see inset below). The walls of the house remained only as wooden frames, all plaster work gone. Gone too was the entire roof, so that the ruin exposed its interior to our lens from all angles, allowing the detailed reconstruction you will find on the following page. This point-cloud may be seen in anaglyph 3D at <http://openendedgroup.com/index.php/software/spatial/bb-3-anaglyph/>



HOUSE WITH TV ANTENNAE, BOMBAY BEACH—MISRACH 1985





Disclosing itself so freely to the lens, this spatialized structure becomes a marvel to explore in 3D, as can be hinted at here with these exterior shots as well as those of the interior shown on the next page.





⑤ *The end is near.* Though just a portion of this structure is visible in Misrach's original shot, it matches up perfectly with the ruin we found there — a structure reduced to just one of its plywood walls, stripped of its original white paint but now covered in graffiti'd pictures and slogans. From one of these we derived its name.

This point-cloud may be seen in anaglyph 3D at <http://openendedgroup.com/index.php/software/spatial/bb2-anaglyph/>.





BOMBAY BEACH: 1. DEAD CORVINA 2. NOOSES 3. GRAFITTI
(IPHONE SNAPSHOTS BY RICHARD MISRACH, DECEMBER 2009)

CONTINUOUS CHANGE

When a site has borne the brunt of radical catastrophic change, one is tempted to regard the now vacated landscape as largely static outside the very gradual weathering processes brought on by sun, wind, and rain.

At Bombay Beach, though, this presupposition is misleading. For one thing, catastrophes are not infrequent on the Salton Sea. Walk toward the water's edge, and your shoes crunch down on the desiccated carcasses of corvina and other kinds of fish layered in the sand, the result of massive marine die-offs that occur on a regular basis.

Even more interesting — for being so much less expected — are the continuous human modifications of the site. State highway 111 runs the 35 mile length of the Sea's western shore, with Bombay Beach the only town to be encountered along the way, half-way down. It attracts a surprising number of onlookers, quite a few of them intent on leaving their marks there.

Such marks come not only in the expected form of graffiti, which are modified and even overwritten with unsettling frequency, to judge from a mere two-month interval in fall 2009. More startling are interventions that we might call *folk art installations*. Everywhere you find little anonymous arrangements of found objects (anonymous for being very rarely tagged — fulfilling purposes distinct from those of graffiti). You might spot a plastic baby-doll sitting upright in the sand, or a green toilet seat propping up a huge plank of plywood, or running shoes dangling from stranded lengths of telephone wire overhead. Misrach once encountered a series of string nooses tied with surprising delicacy to slender trees near the shore, signifying who knows what.

We'll ponder the implications of these interventions in concluding remarks at the end of this paper. Now, however, we proceed to a second, very different Misrach site.



DESERT FIRE #1 — MISRACH 1983

Desert Fire

The Fires section of the *Desert Cantos* was shot during roughly the same time period as *The Flood*.

In contrast to the practice of patience that we saw yielding the Bombay Beach photographs (and most others in Misrach's body of work), this series of photographs from *The Fires* was necessarily the outcome of a fleeting chance encounter — of a fire glimpsed from the window of Misrach's VW bus as he happened to be driving along nearby.

Proposed by Misrach as a second object of our investigation, *Desert Fire #1* first struck us as an even less promising opportunity than the *Submerged Trailer Home*.

We had two doubts. First, we surmised that the blaze would have consumed everything in its path, leaving us almost 20 years later with no landmarks left to match. Second, we didn't see how the site could ever be pinpointed again — to our east-coast eyes, that stand of palm trees looked interchangeable with what we imagined to be countless others dotting the Southern California terrain.

Misrach, however, was perfectly confident that he could lead Kaiser to the exact spot.

| 33



DESERT FIRE #11 — MISRACH 1983



DESERT FIRE #3 — MISRACH 1983



DESERT FIRE #6 — MISRACH 1983




THE THOUSAND PALMS SITE

The site indeed proved easy to find. Some 50 miles north of Bombay Beach, it is a long stand of green trees nestled along the base of some tawny foothills. The nearest house proved to be no more than a mile away, Interstate 10 three miles away, and the larger town of Desert Palms about ten miles away (strip malls, shopping centers, car dealerships, the Motel 6 where we were staying — civilization).

It was early evening when Misrach took Kaiser to verify the spot, and they had to hike a short way across a space that now fell under some sort of protected status. Far from having vanished, the palms there were flourishing, and even in the dim twilight it was a simple matter to pinpoint the precise site of the fire by the trees' charred but living trunks.



Several days later Kaiser returned alone, taking 542 photographs as he approached the site and then partly circled it, scrambling a short way up the choked incline beyond. We had wondered whether the irregular features of the natural landscape would prove more resistant to the spatialization algorithm than the right angles and flat surfaces of built structures, but this proved no problem: the later reconstruction was flawless, and we easily matched the photographs Misrach had originally taken that day.





DESERT FIRE #1 — MISRACH 1983



You may view this point-cloud in anaglyph 3D at
<http://openendedgroup.com/index.php/software/spatial/palm-anaglyph/>.



Misrach had pointed out an oddly curved trunk in the original photograph. Dense foliage now shielded it from view from the same angle, but it was easily made out from the other side.



DESERT FIRE #3 — MISRACH 1983



After the firefighters had put out the fire and had posed for this portrait (a type of photo Misrach rarely takes), they explained the fire's likely cause: teenagers taking target practice on the exposed slope to the right of the trees, where some makeshift targets were spotted. No trace of these remained 16 years later.

DESERT FIRE #6 — MISRACH 1983



Choreography of framing

A possibility we'd raised with Misrach proved as intriguing to him as it was to us. We'd speculated about another way of looking outside the photographic frame — this time for the purpose not of peering at the surrounding space so much as of tracking the photographer as he moved through that space to frame his subject.

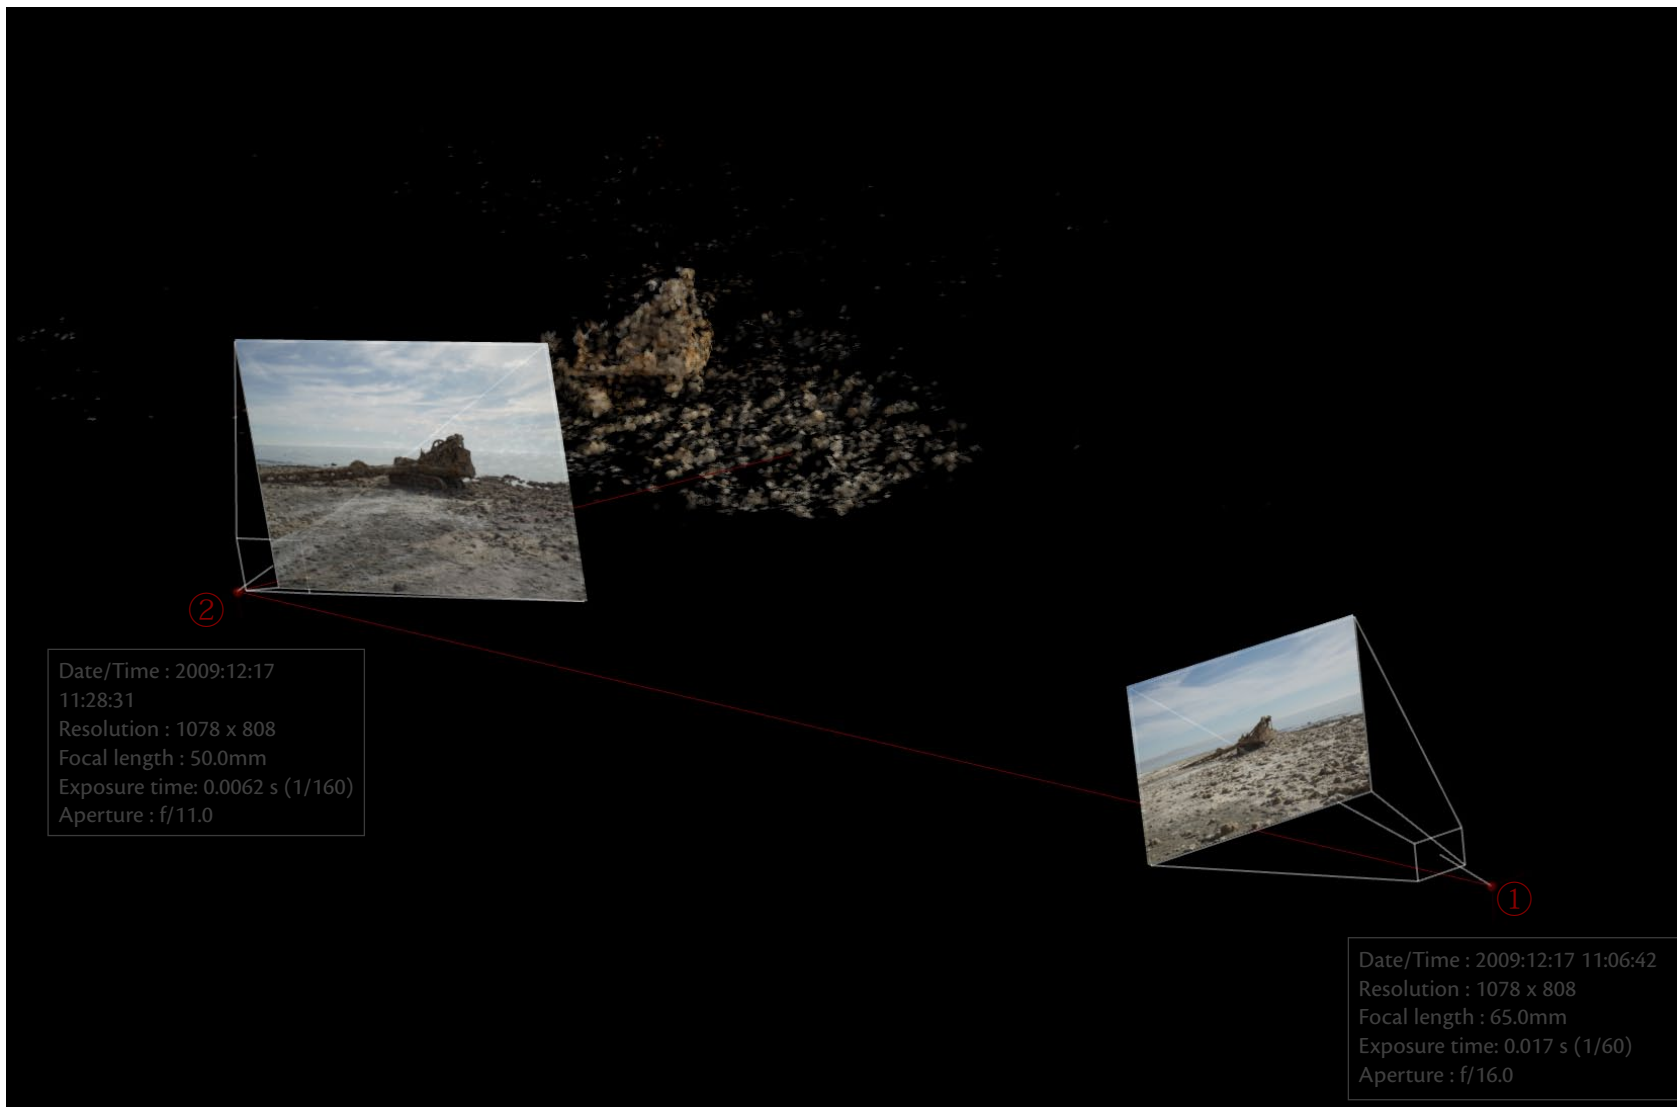
The choreography of framing was an idea that our group (with Shelley Eshkar its third member) had explored in several digital artworks, but these had played out in virtual rather than actual space. The opportunity now was to shadow a real photographer negotiating a real location.

About a quarter mile down the shore from the Bombay Beach site, stranded at the very edge of the water, was a bulldozer. It had at one time been submerged so long that now it had taken on a kind of science-fiction presence, like something out of JG Ballard, all encrusted with barnacles and hardened mold that deformed the lines and surfaces of what had once been the sharp angles of an efficient industrial machine.


This had caught Misrach's eye, and he'd tried shooting it on several occasions, never to his satisfaction. When he returned to Bombay Beach two months after Kaiser was there, he made another attempt at this work-in-progress.

By now he had taken to using a medium format Hasselblad camera body with a Phase digital back, which allowed him to move quickly and shoot often.

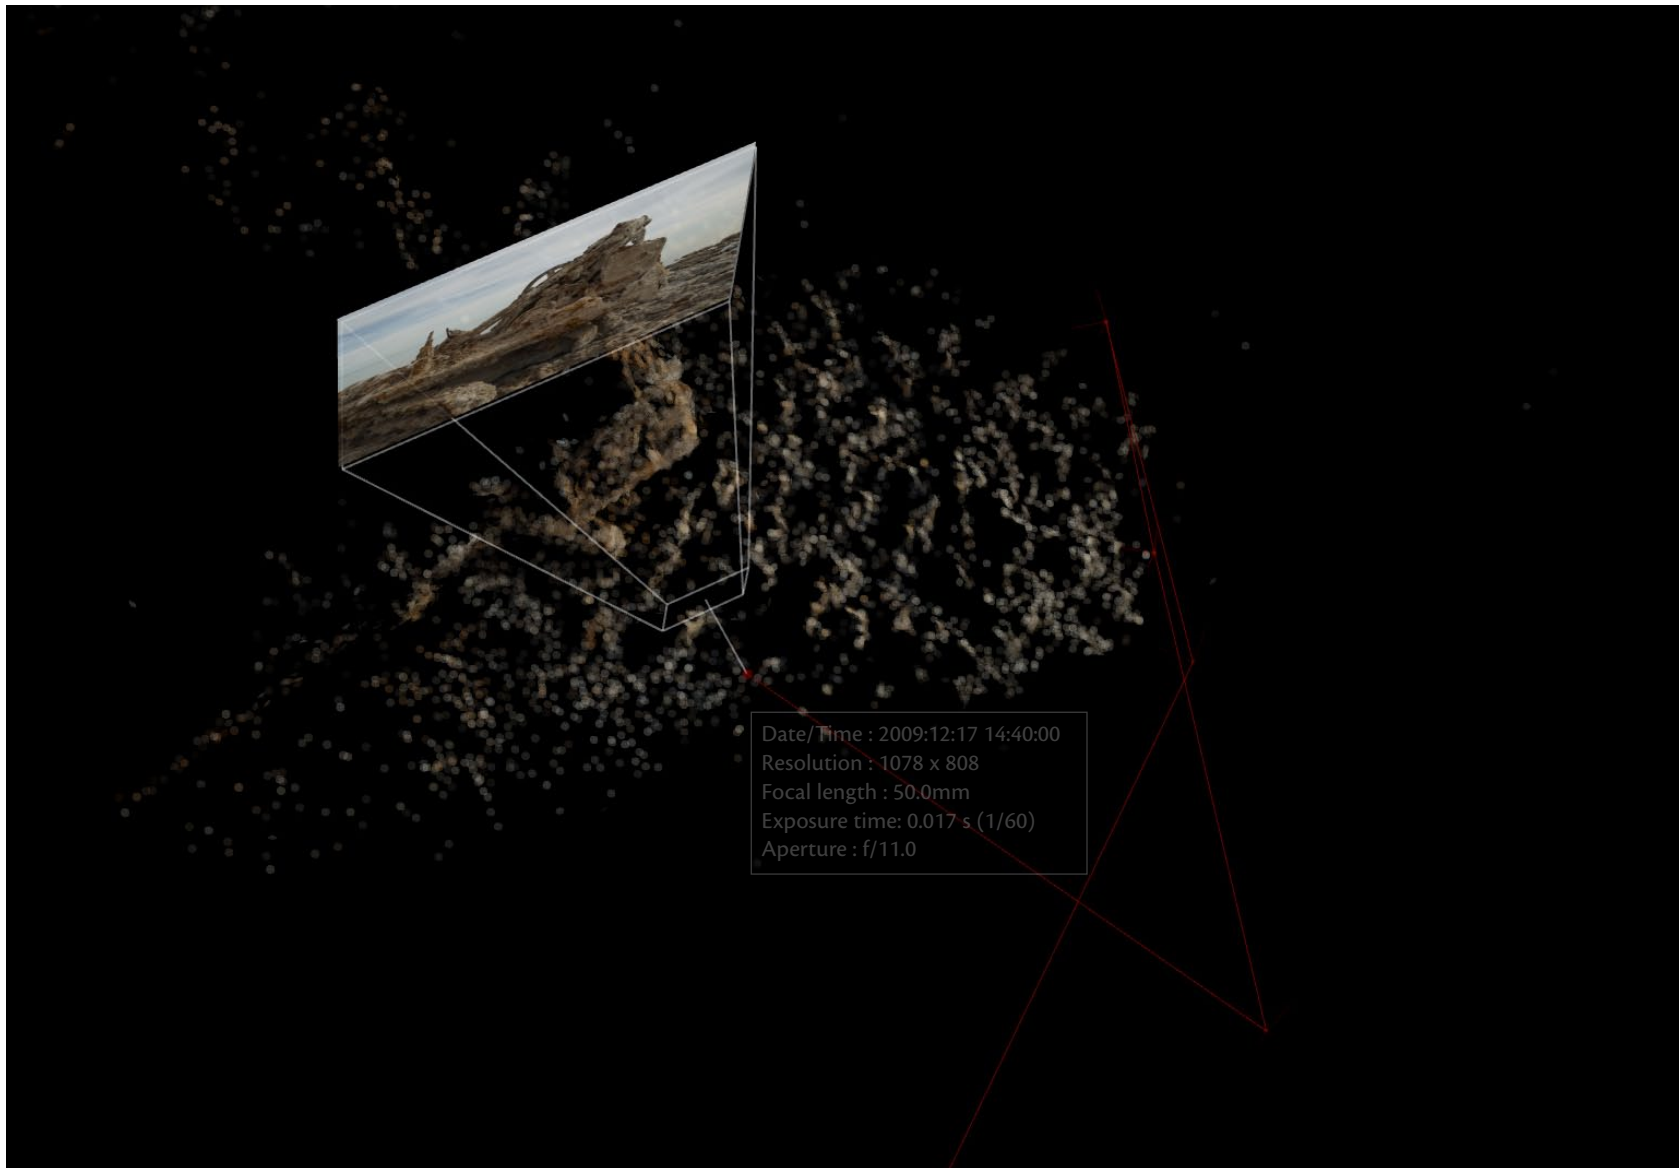
Though in October Kaiser had taken 618 photos of the bulldozer (two of which are reproduced at left), and there were more than enough to generate a detailed point-cloud spatialization, we decided to avoid that approach here. We wondered whether we could spatialize the site and track Misrach's framing of it just on the basis of about 60 photographs that he'd taken there.



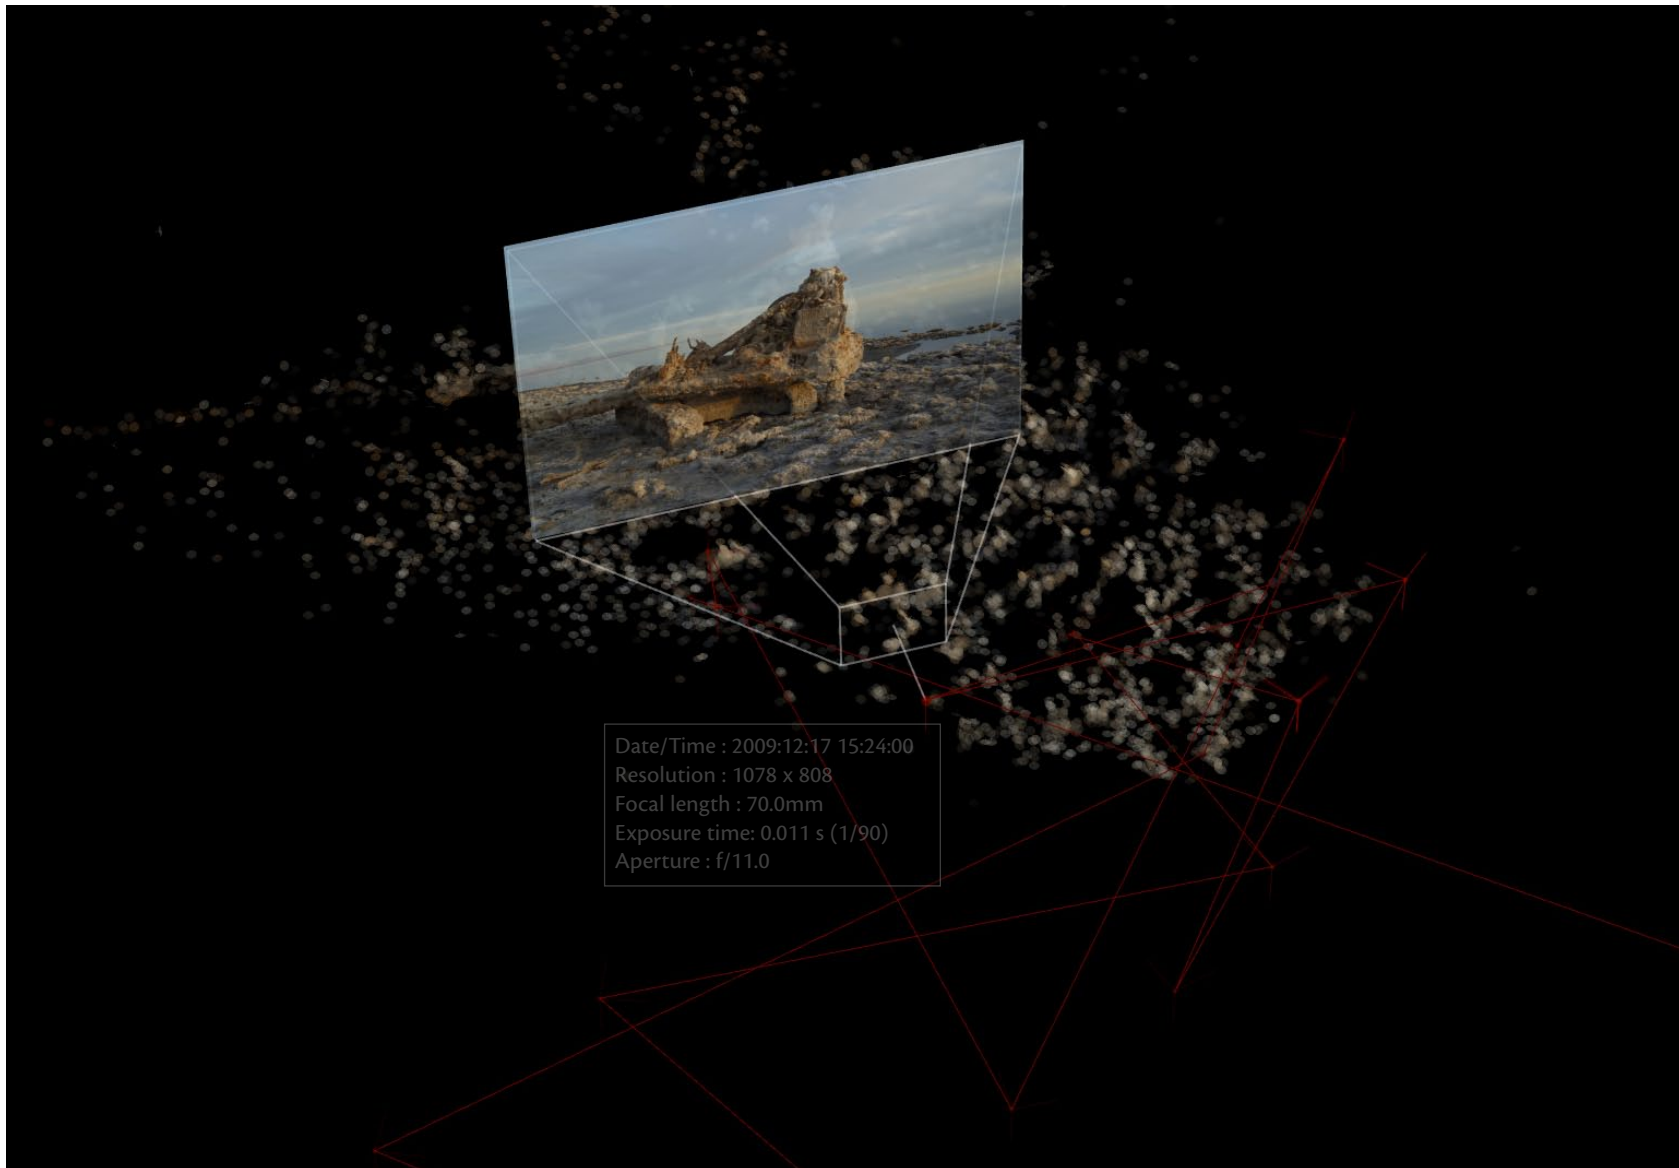
The metadata saved with each digital photograph gives such useful information as the exact date and time it was taken as well as its focal length, exposure time, and aperture. (Note that the resolution given here is misleading: it's for the photos after they were drastically downsampled for our purposes.) The spatialization allows us to pinpoint Misrach's vantage point and angle for each shot, and we can follow the choreography of his approach.



Now the metadata make clear that this third photograph of the bulldozer was taken a little more than three hours after the first two — Misrach had gone off to shoot other subjects in the vicinity and had then returned to this one.



Here we can peer down at the reconstructed site, giving us a better view of Misrach as he comes closer to frame a closer shot.



Date/Time : 2009:12:17 15:24:00
Resolution : 1078 x 808
Focal length : 70.0mm
Exposure time: 0.011 s (1/90)
Aperture : f/11.0

Another skip forward in time: 44 minutes later, and the light and cloud formations are more promising (though in the end not good enough to satisfy Misrach: he fails to get the photo he is after here). We are near the end of his session that day.

You may notice that the camera positions indicated in red exceed the number of bulldozer shots we've documented. This comes from his having turned away to shoot other things: mountains in the distance, birds overhead, dead fish at his feet.

Integration with other data: mapping

One of our original project goals was to interconnect our spatial reconstructions with related information found in other forms of data.

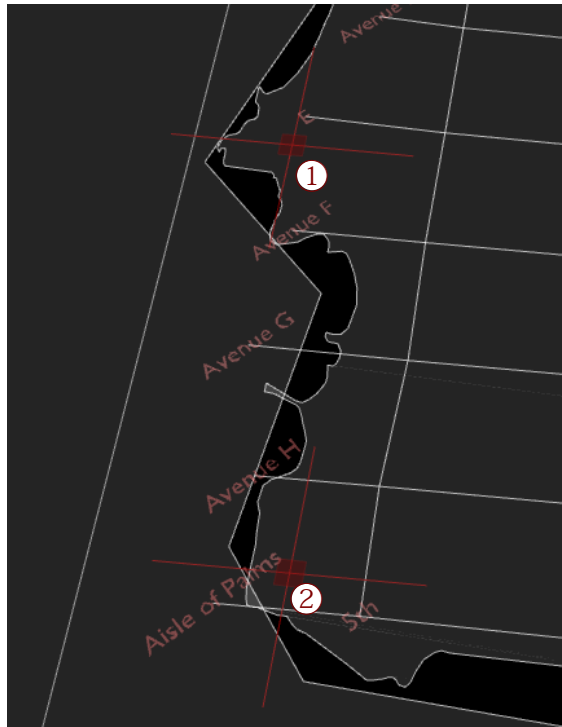
An obvious gateway to data-mining is by means of geotagging, and to that end Kaiser had attached a Nikon GPS instrument to his camera. We soon discovered that comparatively low-cost GPS units of this kind cannot deliver the precision one would imagine and certainly not enough to add any significant information to the spatialization algorithm — giving only five significant digits of latitude and longitude, such units have an accuracy of only plus or minus 65 feet.

However, even these rough geographic coordinates are enough to get us in the game. Recorded as metadata in each photo file, they are the key to unlocking all kinds of web services for performing sophisticated functions with latitude and longitude information — an area that has exploded in the last five years.

For our purposes we chose two such data-bases. The first is OpenStreetMap, an online database that serves up geospatial data. As its name implies, OpenStreetMap is a non-commercial project built up from contributions by volunteers and drawing on other publicly available data-sets. Since what such volunteers select to contribute may differ, what you can retrieve for any given site is likely to vary widely.

The drawing system of our software platform, Field, is flexible enough to retrieve data from the OpenStreetMap server and draw it to the screen instantaneously, as illustrated at left. Bypassing the software tools provided by OpenStreetMap itself gets us much greater flexibility, a point worth examining a bit more closely here.

It is hard to make finished graphical interfaces for extremely heterogeneous databases. OpenStreetMap data consists of three ba-



① Submerged Trailer Home site; ② Bulldozer site.

The camera's GPS data from two of the Bombay Beach locations let us plot their locations with OpenStreetMaps, which serves up the coastline as well as the streets and street-names.



Satellite imagery of Bombay Beach, retrieved from Google Maps.


We note that while Google serves satellite data originally created by NASA, as a commercial service it makes it far easier to retrieve than does that governmental body, whose procedures bear all the marks of a tangled bureaucracy.

sic elements: (1) points in space; (2) paths running through those points); and (3) dictionaries whose entries are tied to either of the two spatial elements. Dictionary entries are linked to a string of text (akin to a definition), and it's these that vary in kind so widely. To give a specific example, look-ups for Bombay Beach will give you the town and street names as you might expect (uploaded by a contributor named Fred), but it will also identify what you probably didn't anticipate: electrical substations and pylons (uploaded by another contributor, Bill).

Because the OpenStreetMap database is so rich and variable, it begs to be explored interactively. You cannot know in advance quite what you will find there, and so you have to try looking for different things. Code-based tools like Field permit this kind of on-the-fly flexibility, which is largely closed off from purely graphical-interface tools, even those provided by the OpenStreet-Map community itself.

The second data-source we drew upon was Google Maps, which provides a very straightforward service for serving up satellite imagery. When queried with an easily crafted (and easily automated) URL, it will serve up tiled satellite images for point on earth at a particular zoom level. Since Field has a general-purpose feature for displaying images in the .png format that Google favors, the tiled images are easily drawn to the screen, as shown at left.

And since 3D is also a general-purpose feature of Field, the combined data can all be viewed interactively in three-space, as shown on the following page.



- Google maps
- GPS from camera
- OpenStreetMap (osm)

While we chose to focus this start-up project on exploring the possibilities and implications to be discovered in Richard Misrach's landscape photography, we were mindful that such an exploration would by no means exhaust related avenues of inquiry that had suggested themselves to us.

Thus we also made a few small forays into other subjects, and two of these yielded insights useful enough to merit presentation here.

Our interest was in situations in which one's materials and/or means were limited. Well-funded and extensively staffed technology projects receive most of the attention, but are by their very nature not the ones on which most scholars toil. Our goal was to suggest new methods of working for scholars who found themselves alone and minimally equipped in the archive, library, or the field.

In the first instance, outfitted with a standard laptop and portable scanner, we spent a day studying an older photographic archive in relatively poor shape — US surveillance and combat photography of Okinawa in 1945.

In the second, we experimented with capturing the kind of object a scholar might find in a museum store-room or out in the field and then might wish to record for future reference and study — a recording that he or she might annotate and then share with others. This time we restricted our field equipment to a laptop and smartphone.

Battlefield Archive — Okinawa, 1945


The photographic archives pertaining to the battle of Okinawa are kept by the National Archives and Records Administration at its facility in College Park, MD.

The idea for investigating these photographs came surrendipitously during a Start-up Project directors' meeting convened at the NEH. On seeing a brief presentation of this project, a historian from UC-Santa Cruz, Alan Christy, wondered if our approach might be applied to the historical landscape of Okinawa, a terrain contested not only in the fierce battle of World War II but also in the conflicting efforts afterward to frame and control its cultural meaning. Christy thought that an interactive spatialization of the island might allow new ways to investigate this topic, particularly if the spatialization allowed one to view photographs of the same area from different time periods.

As we considered how to go about this, we had to rule out the method we'd employed with Misrach's sites — of comprehensively rephotographing a place in order to precisely situate the older photographs taken there. What we soon discovered was that the modern-day landscape of Okinawa bears strikingly little resemblance to the earlier one. So many monuments have been erected as to have obliterated the battlefield beneath them — an ironic outcome of commemoration.

We had to make do, then, with whatever was to be found at the National Archives — which proved to be comparatively little. Not only were only a few of the photographs taken of the same subject from different angles, but also we could plug very few key assumptions into the spatialization algorithm (see explanation at left).

In the end, we were able to spatialize only pairs of photographs, but as can be seen in the following examples, these yielded highly suggestive results.




Analog v. digital photography


It is rarely acknowledged that the effectiveness of the spatialization algorithm relies in part on the metadata it can read from digital files — in particular, the focal length of a lens for any particular photograph. Without such knowledge, you must guess what any given focal length might have been. By trial and error, you can often narrow it down to a level of accuracy that will work. This is easiest, of course, if you know that a single type of lens was used, for at least you then know the range of possible variations.

With the Okinawa archival photographs, we had no such knowledge, and so we had to optimize the algorithm for the most ad hoc case.



It's worth noting that one kind of analog photograph would yield excellent spatializations — stereograms. For these, you could be certain that each paired image shared the same focal length and, more importantly, that both were taken from the same easily guessed distance apart.



Here we attempt to spatialize two photographs of the contoured shoreline towards which boats in the middle ground are heading. This generates a dense 3D mesh reconstruction of the hill, with the other elements of the photograph (water, clouds, boats) barely registering.



Same spatialization, but with our view now rotated around the 3D scene.



A different pair of images yields a spatialization that we show here from two different angles.

Toward a scholar's 3D notebook

If our method requires a collection of numerous overlapping photographs of a subject, there is no reason why that collection should not be assembled from video frames rather than from photos.

We'd realized this early on, and had even considered using a film or video archive as a test case. We had to rule out this idea because we had envisioned processing a massive number of video frames which would have exceeded the computing power that we had on hand for our study.

Later we realized we could work at the opposite extreme — take a very simple test case using video. And so we decided to demonstrate a spatialization workflow that depended only on the video capabilities of a smartphone and the processing power of a laptop computer.

The scenario we imagined was of a scholar studying an object like the small antique clock pictured at left. We thought that this hypothetical scholar would find it useful to capture the object in 3D — not so much as a definitive record, but as a quick entry in a sort of 3D notebook.

Our thought was that the ability to make this kind of 3D capture cheaply and straightforwardly might add a useful tool to a scholar's tool-chest. As an object of study, the capture could easily be shared with others; moreover, the Field workbench could be extended rather simply to enable the scholar to make contextual annotations in the 3D space.

We present the results of our test case on the following page.



FRENCH TRAVELING CLOCK, EARLY 19TH CENTURY




The clock was not the easiest of objects to capture (and chosen for that reason). In addition to its glass surfaces, which at certain angles gave strong reflections (see below), it had an intricate geometry, especially of the gears inside.

We decided to try adding an experimental “point-cloud densifier” capability to our workbench software so that it would render the point-cloud more solidly.

Using an iPhone 4, we shot a somewhat jittery (and certainly inexperienced) video of the clock as we rotated it on a lazy Susan. Selecting every twelfth frame (about 2 frames a second), we fed the stills into the spatialization workbench running on a MacBook Pro laptop.

Ten minutes’ processing yielded the nicely detailed point-cloud to be seen at left. This may be seen in anaglyph 3D at <http://openendedgroup.com/index.php/software/spatial/clock-anaglyph/>.



We set eight goals for this project, all tightly inter-dependent:

1. to refine a method for creating 3D data-sets from unstructured collections of photographs
2. to provide new ways of examining a given point-cloud reconstruction, with any given photograph and the corresponding position of the photographer situated within it
3. to display these visualizations in stereoscopic 3D
4. to make an open-source and flexible workbench for applying this method to advanced scholarly research — a software tool with a dual mode:
 - for novice use of its graphic user interface as is;
 - for expert use at the coding level to freely modify both that interface and its underlying capabilities
5. to tie other data-sources to 3D spatializations
6. to set an example of a new, open-ended approach to the design of scholarly tools
7. to investigate photography from a radically new perspective — by stepping outside its two key defining properties: the framing of a subject and the freezing of a moment
8. to suggest a new form of inter-disciplinary project.

Limitations

Before reflecting on some of the implications arising from our successful realization of these goals, we should first examine two limitations.

The first is that despite our original intention we did not manage to tie a photographic archive to contemporaneous text sources (news articles, eyewitness accounts, and so on). Even with the

Edward Tufte

The best-known advocate for — and exemplar of — visual scholarship is Edward Tufte, whose work has had widespread influence both on the visual analysis and on the presentation of evidence.

Tufte's books are self-published (by his Graphics Press) in beautiful editions that embody the principles he expounds within them. See *The Visual Display of Quantitative Information*, 2nd edition (2001); *Envisioning Information* (1990); *Visual Explanations: Images and Quantities, Evidence and Narrative* (1997); and *Beautiful Evidence* (2006).

We do question, however, one of Tufte's conceptual presuppositions — you'll find our argument summarized in an entry at <http://openendedgroup.com/index.php/2007/09/05/mistaken-manuals-of-style/>. This bears on certain distinctive ideas underlying the approach we are presenting here.

expert assistance of the Aperture Foundation, we could not find a self-contained and open archive both of the right size for this pilot project and with the right set of outward links to related text sources.

What we managed instead was to show how to exploit the GPS coordinates recorded in a photograph's metadata to automatically link a 3D data-set to online geospatial databases. From there, we think it easy to extrapolate ways to link out to much broader text databases. Below we explore a large expansion of our toolbase to permit this kind of data-mining on a massive scale; and we reflect on the kinds of project this would enable.

The second possible limitation arises from a basic working assumption we've made about the emerging present or near-future. In creating an open-source tool whose code researchers can modify and extend freely to match the precise requirements of their investigations, we have hypothesized the emergence of a new kind of humanities scholar — one who, given a suitable digital environment, is capable of meeting the discoveries of computer science and engineering half-way.

For such scholars, we will continue to provide an expanding open-source environment whose capabilities can underpin not only investigations of the kind explored in this study, but also many others. Our hope is that humanities scholars will join scientists and other artists in contributing their expertise in expanding this platform further — and that as these three communities start relying on a common tool-base, they will also forge much richer forms of inter-disciplinary collaboration.

A related assumption is that a genuine *visual* scholarship will develop to match the sophistication and authority of traditional textual scholarship.

Field capabilities documented

You may get both a broader and more specific sense of Field's built-in capabilities by examining the tutorials and the animated walk-throughs to be found at:

<http://openendedgroup.com/field/wiki/StandardLibrary>.

Visualization & data-mining of large data-sets

For an NSF grant entitled *EAGER: Field: A New Tool for Creative, Interdisciplinary Visualization of Data [IIS-1048440]*, we will make Field an advanced tool for the interactive visualization of large datasets on the basis of data-mining. To that end, these three major additions to Field:

1. **Visualization plug-in** — to give Field users quick access to traditional devices for scientific visualization (ie, graphs, charts, and networks) while retaining the option of creating new forms of visualization using the full power of Field's advanced graphics system, including 3D rendering and display.
2. **Data mining plug-ins** — to provide a rich workbench with which to interact with existing data-mining frameworks and libraries, as well as for ones for vast data storage and access.
3. **Linux port** — to enable the use of commodity hardware, thus permitting very large but affordable visualizations on multiple coordinated displays; and to make possible the delegation of complex tasks to GPU-based clusters and ultimately to supercomputers. This will also ease future ports of Field's GUI to other platforms, both large (Windows) and small (Android and iPhone OS).

Field capabilities & ongoing expansion

Our Field software environment already has a very rich set of capabilities that you may draw upon to expand the spatialization workbench (and any other you choose to create). For example, if you have a need to display a graph, the code and interface for making one are there in Field already; or if you want to display a contact sheet of images this capability is also already in place (and the images will automatically thumbnail and load on the fly so as not to paralyze your machine). Field is in constant ongoing development, so even when the spatialization workbench is not the current focus of our coding, its capabilities are automatically improved with any expansions of the over-all environment.

Such expansions are most often generated by the diverse requirements of the digital artworks we create. However, our next major enlargement of Field will arise from a National Science Foundation grant, for which we will greatly strengthen Field's data-mining and visualization capabilities, enabling it to interact with very large data-sets (see summary at left). Conducted jointly with scientists and scholars at Rensselaer Polytechnic Institute, the year-long project commences in January 2011. The resulting expansions to Field's codebase will be released as they are made, and so will be promptly available for you to use.

Advantages of scalability`

These added data-mining and visualization capabilities will enable very large and ambitious projects — but for now let us return to considering the opposite.

With our idea for a scholar's 3D workbench (p. 52), we outlined the kind of small lightweight tool that by making use of everyday digital equipment allows the typical scholar to conduct new kinds of fieldwork, analysis, interchange, and presentation.

It's instructive to compare this with the nearest existing equivalent: the use of a laser range-finder to build a very accurate 3D point-cloud of a given structure or space. The advantage of this approach (its accuracy) are offset by its very large cost of deployment. Requiring an expensive technical set-up operated by a large team of trained technicians, it entails the systematic capture of the 3D data over a relatively long period of time.

And because of its high cost in time, money, and expertise, laser range-finding technology is typically applied to well-established subjects, whose high value is known in advance. Thus, typical laser range-finding projects include exhaustive captures of Michelangelo's David and of ancient Buddhist temples.

What this rules out is the low-cost, ad hoc, experimental investigation, in which the value of the outcome to the world at large may be uncertain, but which is certainly of keen interest to the individual scholar pursuing it.

For this reason, it is crucial to put tools like the visualization workbench into the hands of scholars who do not necessarily command very large research budgets and who may choose to investigate less prominent historical or cultural artifacts.

Since Field scales well, it allows the addition of resources to its pipeline as they become available. Thus a scholar who starts by working only with a laptop in the field may, on returning to the office, add a networked desktop computer or two to the spatial processing tasks.

If the project then increases in complexity and could use greater processing power, users will soon be able to add not only other Macintoshes to the pipeline but also commodity hardware running Linux and even GPU-based supercomputers.

To see what projects might warrant this kind of scaling up, we will conclude by exploring two hypothetical ones.

Hypothetical example of data-mining — Ground Zero

The broadest opportunities for using 3D photo spatialization clearly lie in reconstructions of the recent past — for the obvious reason that the number of images taken of any given event is now vastly greater than ever before. Moreover, since most of these images are born digital, many are accessible online and all carry metadata, which as we have seen aids techniques both of 3D reconstruction and of data-mining.

Since the difficulty we all face now arises not from a paucity of images and information but from a gushing surfeit of both, we need to forge ways not only of organizing this overwhelming profusion, but also of making new sense of it. Combining data-mining with photo spatialization suggest one possibility for doing so.

We can take as our hypothetical example the 9/11 destruction of the Twin Towers. With this disaster having taken place in the media capital of the world and occurring not in an instant but over a span of two hours, no event has ever been more thoroughly recorded — from an extraordinary diversity of vantage points and types: helicopter broadcast feeds, cellphone snapshots, news photography, amateur video, etc. Add to this non-visual sources such as news reports, text messages, voice and email, blogs, eyewitness oral histories, radio dispatches, police logs, and so on, and you have intricate intertwining streams of visual and textual information.

9/11 at Ground Zero has yielded itself fairly easily to timeline reconstructions and associated narrative histories — that is, to traditional forms of textual scholarship. The visual materials have also lent themselves straightforwardly enough to documentary film, but this, we would argue, takes the same linear form.

What we have lacked is a way to make sense of the overlapping and overwhelming simultaneity that actually prevailed there —

of the bewildering number of viewpoints brought to bear on just about every second of the event as it unfolded.

We can now imagine being able to navigate the vast three-dimensional space of a single second.

We could move from a view of the burning towers framed by an apartment window to one with the fire looming right overhead to another looking down on the structures from the air to yet another looking out from a window in one of the Towers itself. Correlated with this moment would be what one might call *non-visual viewpoints*: a voicemail, a text message, a police log notation, etc, all with the same timestamp, and many geotagged and therefore also situated spatially.

Thus by placing simultaneous viewpoints in a coherent spatialization (itself built in part from those views), we can establish a visual way of knowing quite distinct from the linear textual one.

| 59

Hypothetical example of interdisciplinarity — Bombay Beach

We may extrapolate our second example from the Bombay Beach site presented earlier. Though also a site of devastation, it differs markedly from the Twin Towers example — for where the histories of Ground Zero are public, prominent, and exhaustively documented, those of Bombay Beach are private, obscure, and sparsely (but interestingly) captured.

A project on Bombay Beach would require a concentrated act of excavation and recovery — an act that would have to cross disciplines broadly if it were to interconnect all the crucial elements.

For we'd need to try reconstructing: (1) the personal histories of the former inhabitants; (2) the general history of the settlement; (3) the prior as well as the ongoing photographic record of the

place; (4) the ever-changing graffiti together with the even more transient “folk art installations” (as we’d termed them earlier); (5) the unexpected dissemination of Bombay Beach imagery to distant locations; and (6) the past, present, and projected ecological condition of the site.

All of these elements may be interconnected spatially. For a particular spot on the site, we may find a life that was lived there, a photograph taken, some graffiti tagged, a flood line breached, an insurance claim denied, and so on.

We would argue that such a project must cut across the boundaries commonly — and even officially, in the charters of the relevant federal agencies — established between the humanities, the sciences, and the arts. For the research would require a collaborative team with expertise in oral and cultural history, sociology, ecology, computer science, photography, and immersive installation art.

Here we will devote the remainder of our examination to aspects of the project too often roped off as art.

First, the art of photography — for Misrach’s original photographs of the site put Bombay Beach on the cultural map. Not only did his pictures of the flood there receive very wide exposure and become part of the art photography canon, but the kind of approach he and kindred photographers took to the western landscape have diffused into the culture at large, often spreading so far that the influence is exerted without being named or known.

By this we mean that, as we discovered, a taste for photographing the Bombay Beach ruins was surprisingly prevalent. In our time there, we noticed quite a few people pulling off the highway to photograph the place. Misrach even encountered a Japanese photographer there who knew of the site from having seen *The Flood* photographs and had come to America only to photograph



This photograph, of a ruin we recognized as being a stone's throw away from the site of Misrach's original photos, was designated an honorable mention in a photo competition we found displayed on the outer wall of the Hearst building in New York, with a label identifying it as *Salton Sea Trailer*, 2008, by a David Zimmerman.

it himself. And back in New York, we were surprised to stumble upon Bombay Beach even on Broadway, as is illustrated and recounted at right. Later we came across a book by Kim Stringfellow entitled *Greetings from the Salton Sea: Folly and Intervention in the Southern California Landscape, 1905-2005* (Center for American Places, 2005), with pictures bearing the same Misrach DNA.

But without question many amateur photographers are drawn to ruins like Bombay Beach not because they have seen the original Misrach photographs (or, for that matter, have read the dystopian fictions of JG Ballard), but because the aesthetic tradition these artists have advanced has now seeped out into the general culture.

Even more surprising is how widely the aesthetic of the found object and appropriated space has also propagated. Robert Smithson's great 1967 essay "A tour of the monuments of Passaic, New Jersey," first published in *Artforum*, may have been read by very few, but its sensibility has taken root in more places than one might imagine and now frames many people's views of marginal territories like the Salton Sea.

Smithson is of course best known as an earthworks artist — and earthworks may now serve us as a telling point of contrast. Enormous interventions like Smithson's *Spiral Jetty*, which bulldozed part of the Western landscape into a monument, are on a scale akin to Egyptian or Mayan pyramids. But rather than the collective expression of a whole culture like those ancient works, the *Spiral Jetty* is the vision of a single man, and from a certain distance may seem to border on megalomania.

If we can imagine creating Bombay Beach as a human-size immersive 3D installation, it will have something in common with Smithson in its sense of scale, of entropy, and of inexorable natural forces. However, it would be peopled not only by the artist or artists who created it, but also by many others who have used the same canvas of

Bombay Beach as their means of expression — from the graffiti tags with which they imprint the place, to the found art arrangements they leave behind, to the photographs they post online.

And many of the stories told there would be by those whose memories situate themselves within a vanished space whose vestiges have been reconstructed in this simulacrum. Such a virtual installation, we believe, would be intrinsically hybrid — an artwork that enfolds many art expressions; a history that intertwines personal, cultural, historical, and ecological strands.

We could not have created this study without the unstinting help of **Richard Misrach**, who took Kaiser on a detailed tour of four key sites at which he took some of the now-classic photographs of his lifework, the *Desert Cantos*. (The pressures of time kept two of these sites out of our study, though they were equally rich in implication.) Richard not only shared unpublished workprints with us, but also went back to the Bombay Beach site on his own to rephotograph it (which allowed us to create the “Choreography of framing shots” section). During his photographic journeys, he kept us up to date with startlingly good snapshots taken with his iPhone. Richard also made the resources of his studio available to us; his assistant, **Seza Bali**, kindly retrieved and scanned the workprints that we relied upon.

Melissa Harris, editor of *Aperture* magazine, encouraged us to create this project in the first place and then helped us consider possible photographic archives for study; it was she who introduced us to Richard Misrach. Together with her colleague **Leslie Martin**, publisher of Aperture Books, Melissa pointed out historical antecedents for us, drawing our attention especially to the remarkable work of the Rephotographic Survey Project.

Our exploration of the Okinawa battlefield archive was suggested by **Alan Christy**, whose UC-Santa Cruz project on Pacific Memories had found support in a concurrent Digital Humanities start-up grant. His graduate student, **Amanda Shuman**, kindly conducted an advance survey of the Okinawa archives at the National Archives.

Though not officially part of this project, **Shelley Eshkar**, our long-time colleague in OpenEnded Group, advised us in his reliably insightful fashion. It was his suggestion that by exhaustively re-photographing a given site, we’d be able to situate related archival photographs, often too few in number to support a spatialization on their own.

We are deeply indebted to all those programmers whose work made ours possible.

Specifically, our visualization workbench packages these open-source projects:

Bundler — Structure from Motion for Unordered Image Collections. Noah Snavely, Steven M. Seitz, Richard Szeliski. “Photo Tourism: Exploring image collections in 3D.” *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2006)*.

PMVS2 — Patch based multi-view stereo by Yasutaka Furukawa and Jean Ponce. “Accurate, Dense, and Robust Multi-View Stereopsis.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2009.

sift++ — a lightweight C++ implementation of SIFT detector and descriptor. Andrea Vedaldi.

PrimitiveShapes — R. Schnabel, R. Wahl, R. Klein — “Efficient RANSAC for Point-Cloud Shape Detection,” in *Computer Graphics Forum*, Vol. 26, No. 2, pages 214-226, Blackwell Publishing, June 2007.

CGAL — Computational Geometry Algorithms Library, <http://www.cgal.org>.